

Internet Firewalls Policy Development and Technology Choices

Leonard J. D'Alotto
GTE Laboratories, Incorporated

Abstract

Since the development of the World Wide Web (WWW), more and more organizations are connecting their networks to the Internet. Many of these organizations are, rightly, concerned about the security of these connections. Realizing this, a number of companies are producing products known as Internet Firewalls and marketing them as the security solution. Faced with a blizzard of offerings, and a growing "feature war" between firewall vendors, confusion tends to be the order of the day. So what is to be done?

This paper addresses this question by arguing for the need for a proper Internet security policy. The information which should be included in that policy, and ways to use that policy for determining the appropriate firewall technology are given. This paper is based on the author's significant experience in evaluating firewall products and implementing them in a variety of environments.

1. Introduction

Since the development of the World Wide Web (WWW), more and more organizations are connecting their networks to the Internet. Many of these organizations are, rightly, concerned about the security of these connections. Realizing this, a number of companies are producing products known as Internet Firewalls and marketing them as the security solution. Faced with a blizzard of offerings, and a growing "feature war" between firewall vendors, confusion tends to be the order of the day. So what is to be done?

What must be realized is that firewall products are only a part of any Internet security solution, and the part they play differs not only between products, but between organizations and their Internet connections. These roles, the type of firewall product, and the need for features, depends upon the organization's policy towards Internet connectivity. That policy addresses not only security, but use of the Internet by employees and other issues related to the Internet. In addition, an organization must consider how that firewall product will be managed as part of the organizations overall network. This paper addresses the development of Internet policies, looks at available firewall technologies, and provides guidelines for applying the available technologies to meeting that policy.

2. Internet Connectivity Policy Development

Before an organization connects to the Internet, a policy governing that connection should be established. This policy should address three major areas: security, use, and management and administration. The organization must be aware, however, that policy development is a process that is not complete until all firewall and other technologies to be correctly implemented are chosen. In the process of implementing an Internet connection, an organization will usually tighten (or loosen) the policy based on risk mitigation vs. investment.

2.1 Security

This portion of the policy addresses what traffic is allowed to flow between the organization's network and the Internet. In setting these policies for an organization, more than just the security risk of a connection must be considered. For some organizations, inbound telnet is a real requirement. Inbound telnet is dangerous. Depending on the organizations mission and financial strength, strong user authentication and possibly encrypted sessions may be required. The key is to determine, for each potential Internet service:

1. Is there a real requirement to allow this service, both inbound and outbound,
2. What are the risks to the organization from this service,
3. What is the level of investment the organization is willing to make to mitigate these risks,
4. What are the preferred mitigation methods, and
5. What is the method for handling new Internet services?

Upon compiling this information, an organization can then proceed to develop the policy on who has access to Internet services.

The one item in this section that is frequently overlooked is the issue of how the organization will determine if new Internet services will be allowed. As new Internet services are developed, individuals within an organization will request access to these services. Since many of these services require 2-way communications, they could potentially cause an opening in the firewall. Therefore, a policy and set of procedures on how to request these services, and how to decide whether access to them will be allowed, must be developed. Otherwise, there will be no way of properly managing these requests and maintaining control over the firewall.

2.2 Use

This section of the policy addresses which members of the organization will be given Internet access, and the type of access they will be given. It is not unusual for an organization to give all employees the same access to the Internet. On the other hand, many organizations limit access to a "select" group of individuals, and then further restrict services within that group. For example, all professional staff may be given E-Mail access but only marketing and research (and of course, all top executives) may be given access to the World Wide Web. The decisions made in

developing this section of the policy are crucial to determining the type of firewall technology to be used.

The second part of this section involves determining the way in which access is to be restricted. Is it to be based on IP address, user ID and password, or via strong authentication of the user? All three options are available, but their impact needs to be understood. If access is to be limited based on the user's IP address, it is easily subverted. This can be done by any individual walking up to the authorized workstation, sitting down, and going to work. It can also be thwarted through IP spoofing and related techniques. The other problem with this is that every time a legitimate user moves, and thereby receives a new IP address, the rules in the firewall must change. With an account based restriction, where the user must authenticate their identity to the firewall, this administrative burden is removed as once an account is established, it is valid from all IP addresses within the enterprise.

2.3 Management and Administration

This section presents the guidelines for managing and administering the Internet connection. In setting these guidelines, there are several questions that must be answered. They include;

1. What events are to be logged,
2. Where are logs to be kept, and for how long
3. What events are to be alarmed,
4. What types of alarms are to be required, e.g., E-Mail, pager dialing, etc., and
5. What types of reports are to be required.

2.3 Remaining Issues

As one can see, by going through this process and documenting these policy decisions, the requirements for the firewall are almost complete. There are two major items missing, however. One is the skills and capabilities of those who will manage the firewall. If your organization is not staffed with capable Unix and TCP/IP administrators, this must be taken into account. One example is an old mainframe and SNA shop. While the people were bright and very professional, they were not expert Unix and TCP/IP administrators. Consequently, when one vendor's firewall was found to require the administrators to configure the system by editing Bourne shell scripts, it had to be rejected. This was in spite of the fact that at the time, it was one of the more robust firewalls available.

Another item to be addressed, and this is closely tied into the above, is who is to administer the firewall. Is this to be in-house, outsourced, or some combination of the two? All options are available in the market, and after careful consideration of the policy and the capabilities of available staff, a proper decision may be made.

Finally, the issue of compliance needs to be addressed. Unless a mechanism is put in place to ensure compliance, then this entire policy will be meaningless.

3. Firewall Technologies

What exactly is a firewall? A firewall is, simply, “A set of tools used to implement an Internet security policy” (heard by a participant in a late night discussion over beers at the November 1995 CSI conference). What this means is that the firewall product purchased from a vendor is not the entire firewall system. Rather, the firewall system includes the use, if applicable, of a protected mail host, a “split” Domain Name Service (DNS), and the use of a DMZ. While we give a brief introduction to these technologies here, the reader is referred to [2] for more information.

3.1 Proxy Servers and Packet Filters

These are the technologies encompassed in a typical firewall product. Simply put, a packet filter restricts based on source and destination IP address and TCP or UDP service port. A proxy server is “...a program that deals with external servers on behalf of internal clients. Proxy clients talk to proxy servers, which relay approved client requests on to real servers, and relay answers back to clients.” [2]. In practice, a proxy server requires users to log into the firewall, and then access the Internet from that server. There are now firewall products appearing on the market which perform packet filtering but require users to authenticate themselves. Once the authentication is complete, the firewall will allow traffic for that session to pass.

3.2 Use of a “DMZ”

A DMZ, a term stolen from [3], is simply a network between the firewall and the Internet over which you have some control. Figure 1 depicts such a network.




Figure 1 DMZ

The purpose of the DMZ is to have a place outside the firewall, but over which you still regain control, for external DNS servers, external mail hosts, public information servers, etc. This way, inbound traffic can be more tightly controlled, and information about the internal network kept from being published to the Internet. For more on DNS, Mail, WWW server, etc., security see [1], [2], [3], [4] and [5].

3.3 Encrypted Tunnels and Virtual Private Networks

The latest trend in firewall offerings is to add cryptographic services for firewall to firewall encryption. The encrypted traffic between these firewalls is referred to as an encrypted tunnel or a virtual private network (VPN). The purpose of this technology is to allow two sites connected to the Internet to use cryptography to communicate in secret and with total security. This then allows a company with multiple sites to use the Internet as a Wide Area Network (WAN) instead of paying for expensive private leased lines.

A word of caution is in order. If one wishes to use a VPN for this purpose, one must make sure the endpoints (firewalls or routers) are secure. The best cryptography in the world is easily subverted should the endpoints be easily penetrated and the cleartext visible to that penetrator.

4. Application of Technologies to Policy

The following case studies are based on actual situations. They are intended to illustrate the types of situations in which packet filtering, proxy servers, and packet filters with user authentication are appropriate.

4.1 Case Study - Packet Filtering

In this particular instance, the organization was currently connected to the Internet. The policy in place, and to be kept, was that all employees are to be given E-Mail, news, and outbound telnet, ftp, gopher, and http. Inbound services were to be limited to mail and news. In addition, DNS needed support, as did an anonymous ftp server and WWW home page. At the time, packet filters on the Internet router were being used to provide a rudimentary level of security. This was deemed inadequate as the rule base to implement a policy was complicated, and proper logging could not be supported on the router.

In this case, there was no need for authentication in either direction. Inbound traffic, being just E-Mail and News, could not be subject to any user authentication. Outbound was to be allowed, regardless of the workstation or user that originated that traffic. Additionally, all traffic could be restricted based on TCP or UDP service port, in conjunction with the source and destination address. (For example, an inbound connection on port 25 to the mailhost would be allowed, but to any other host in the organization, it would be disallowed.)

The result was a packet filtering firewall with a DMZ, split DNS, protected mail host, and appropriate event logging. Since internal user authentication was not required, as everyone had full Internet access, and no inbound traffic other than mail or news was allowed, no user authentication was required. With the split DNS and protected mail host, full service could be offered to users without publishing information as to the structure of the enterprise network. And, since most packet filtering firewalls provide fairly extensive logging facilities, the Internet security policy could be implemented for this organization fairly simply.

4.2 Case Study - Proxy Servers

In a second case, the policy was much more complicated than the first. The organization was extremely large, and people move offices fairly regularly. In addition, many users were on a LAN with a proprietary LAN OS which does not support native IP. This LAN also used a proprietary E-Mail transport with an SMTP gateway. The policy requirements resulted in five types of users - no access, E-Mail only, E-Mail plus news, full access restricted by day of the week and time of day, and unrestricted access. With the various types of access to be given to individuals within the organization, each request for outbound access would have to authenticate the user. With the constant changing of users workstation addresses, address authentication would be impossible to administer. Consequently, this was a classic case of where one uses a proxy server.

In addition, a robust DMZ for a protected mail host and support of a split DNS was required in this situation. The first reason for this is the E-Mail situation. With a proprietary internal mail system that utilized SMTP gateways, there needed to be some way of managing traffic to those gateways. An external mail host was used to simply relay mail to one of several internal mail hosts, on which an SMTP gateway resided. This reduced the processing load on the firewall, as it was not required to determine which mail host a message should be relayed to. E-Mail addresses were set up to be of the form `user@mailhost.domain`, and the external mailhost was configured to only know the addresses of the internal mailhosts. Coordinating this with a split-DNS allowed for separation of the Internet and the LAN-OS based users, as well as protecting against publicizing the structure of the IP based portion of the network.

4.3 Case Study - Filters with User Authentication

Where does one use a packet filter with user authentication? As it so happens, this was not done in linking to the Internet, but rather, in linking contractor and vendor networks to the corporate LAN. In this particular case, a variety of contractors and vendors need access to the organization's LAN, but only to limited machines. The contractors and vendors all came from registered Class B or Class C networks. However, traffic flow from a contractor needs to be limited to those machines to which they are authorized access. This allows for a simple rule base where the rules are of the form:

From	To	Service	Action
Class B or C address	Host Addresses	Service Ports	Allow

The problem is, not all personnel from the contractor are allowed access. Therefore, user authentication is needed. Now, the question arises, why not just use a proxy server? The answer is that the contractors are doing development and maintenance and need a variety of services for which there are no available proxy servers. So, setting up a packet filter with user authentication presented itself as the only alternative. That is, if the traffic matches an allow rule, authenticate the user as one authorized to generate this traffic, and pass it. Otherwise, drop it.

5. Summation

In closing, one can see that preparing a robust policy is a prerequisite before making a choice on how to firewall an Internet connection. If such a policy is written, and the approach to implementing the firewall chosen and documented, these can then be given to firewall vendors. The vendors should then be asked to provide, in writing, a document describing how they will implement your policy in their product. Upon receiving these documents, choose two or three finalists, and have them provide evaluation systems. Place these systems in a laboratory in which the real connection can be simulated. Extensively test these products before making a decision. In this way you can gain a proper understanding of how the firewall product will fit into your organization. From this, you can choose a vendor and implement your firewall.

6. References

[1] Paul Albitz, Cricket Liu; DNS and Bind; ; O'Reilly & Associates © 1992

[D. Brent Chapman and Elizabeth D. Zwicky; Building Internet Firewalls, O'Reilly & Associates © 1995

[3] William R. Cheswick, Steven M. Bellovin; Firewalls and Internet Security - Repelling the Wily Hacker; Addison-Wesley Professional Computing Series, Addison-Wesley Publishing Company, © 1994

[4] Bryan Costales, Eric Allman, Neil Rickert; sendmail; O'Reilly & Associates © 1993

[5] Cricket Liu, Jerry Peek, Russ Jones, Bryan Buus, Adrian Nye; Managing Internet Information Services; O'Reilly & Associates © 1994